URBAN COMBAT DATA MINING

Barry A. Bodt*, Eric G. Heilman, Richard C. Kaste, Janet F. O'May
U.S. Army Research Laboratory (ARL)
Aberdeen Proving Ground, MD, 21005-5067

ABSTRACT

We describe an approach and its implementation involving simulation and data mining for improved understanding of the potential relationships among battle parameters and battle outcomes in an urban setting.

1. INTRODUCTION

The Future Force will rely on new concepts of networks, information, and integration (NII) supported by prolific data collection assets. Novel methodologies will harness the data-rich battlespace to improve the planning/re-planning cycle. ARL has been exploring the applicability of combat simulation to NII, with endeavors centering on course of action (COA) parameterization and assessment. The scope is the development of a consistent set of tools that assist the warfighter in situation assessment, planning, wargaming, and execution monitoring.

Advances in simulation and data mining are increasingly applicable to battlespace decision support. In particular, simulations now generate realistic behaviors at increasingly finer scales, and data mining mechanisms uncover useful patterns in large data sets. Moreover, these advances are facilitated by high performance computing. Data mining detailed COAs in urban scenarios has the potential to yield insights for commanders by providing a capability to examine simulation data for relationships between the COA structure and operational characteristics.

Previous investigation suggests good potential for relating detailed and futuristic battle parameters (e.g., real time entity health and munitions quantities) to battle outcome in a setting involving engagement of a few armor platoons in open terrain (e.g. Bodt et al., 2002, Liao et al., 2003). For that scenario, it was possible to successfully predict battle outcome in a battle averaging ~ 45 minutes as a function of battle parameters measured far earlier in the conflict: 70% correct classification at 5.5 minutes and 85% correct classification at 20 minutes. More importantly, the parameters supporting the correct classification were identified as key parameters for both planning and execution monitoring.

In the present paper we address the more difficult situation of Military Operations on Urban Terrain (MOUT) with armor and dismounted infantry (DI) in a phased attack on an urban location. Multiple objectives, team interaction and a growing number of battlefield entities to track provide a far greater challenge for the approach. In what follows we describe the scenario for our investigation, identify the battle parameters considered, describe the analytical approach taken, discuss the results, and recommend further work.

2. SCENARIO

A sector of a city, based on the McKenna MOUT site, Fort Benning, GA, serves as the terrain. The scenario consists of an attack on a city sector carried out in two distinct phases. Phase 1 consists of the attack to isolate the area. Swift movement characterizes the attack, comprising a three-pronged encirclement to reduce threat forces from positions around the sector. See Figure 1. The first phase is carried out by five Bradley Infantry Fighting Vehicles (M2s) and two Abrams tanks (M1A1s). Two M2s attack from the north; two M2s attack from the southwest and two M1A1s and one M2 (signifying the Headquarters [HQ] attachment) attack from the west. The initial resistance is provided by three infantry fighting vehicles (BMPs) and two tanks (T-80s) positioned immediately inside the perimeter of the sector. A third T-80 is in the center of the sector flanking the critical Objective building. The armored vehicles entering the sector are to stop at the edge and serve as protection for DI in the second phase.

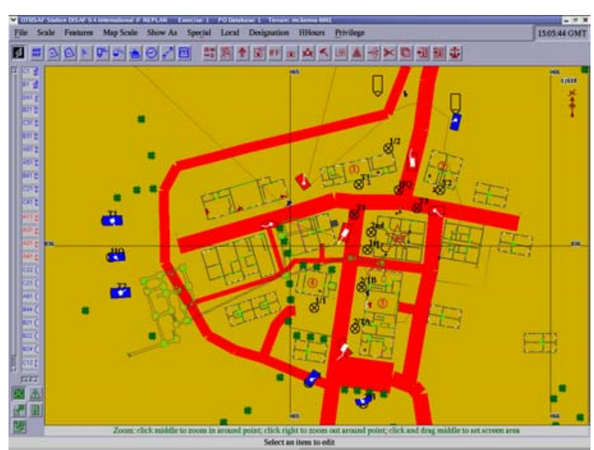

Fig. 1. Encircled city sector in the initial phase

<table>
<tr><td colspan="6">Report Documentation Page</td><td colspan="2">Form Approved
OMB No. 0704-0188</td></tr>
<tr><td colspan="8">Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</td></tr>
<tr><td colspan="3">1. REPORT DATE
00 DEC 2004</td><td colspan="3">2. REPORT TYPE
N/A</td><td colspan="2">3. DATES COVERED
-</td></tr>
<tr><td colspan="6" rowspan="3">4. TITLE AND SUBTITLE
Urban Combat Data Mining</td><td colspan="2">5a. CONTRACT NUMBER</td></tr>
<tr><td colspan="2">5b. GRANT NUMBER</td></tr>
<tr><td colspan="2">5c. PROGRAM ELEMENT NUMBER</td></tr>
<tr><td colspan="6" rowspan="3">6. AUTHOR(S)</td><td colspan="2">5d. PROJECT NUMBER</td></tr>
<tr><td colspan="2">5e. TASK NUMBER</td></tr>
<tr><td colspan="2">5f. WORK UNIT NUMBER</td></tr>
<tr><td colspan="6">7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
U.S. Army Research Laboratory (ARL) Aberdeen Proving Ground, MD, 21005-5067</td><td colspan="2">8. PERFORMING ORGANIZATION REPORT NUMBER</td></tr>
<tr><td colspan="6" rowspan="2">9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)</td><td colspan="2">10. SPONSOR/MONITOR’S ACRONYM(S)</td></tr>
<tr><td colspan="2">11. SPONSOR/MONITOR’S REPORT NUMBER(S)</td></tr>
<tr><td colspan="8">12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited</td></tr>
<tr><td colspan="8">13. SUPPLEMENTARY NOTES
See also ADM001736, Proceedings for the Army Science Conference (24th) Held on 29 November - 2 December 2005 in Orlando, Florida. , The original document contains color images.</td></tr>
<tr><td colspan="8">14. ABSTRACT</td></tr>
<tr><td colspan="8">15. SUBJECT TERMS</td></tr>
<tr><td colspan="4">16. SECURITY CLASSIFICATION OF:</td><td rowspan="2">17. LIMITATION OF ABSTRACT
UU</td><td rowspan="2">18. NUMBER OF PAGES
7</td><td colspan="2" rowspan="2">19a. NAME OF RESPONSIBLE PERSON</td></tr>
<tr><td>a. REPORT
unclassified</td><td>b. ABSTRACT
unclassified</td><td colspan="2">c. THIS PAGE
unclassified</td></tr>
</table>

Phase 2 is highlighted by eight fireteams entering behind the armored vehicles. Three fireteams enter from both the north and the southwest with two fireteams entering from the west. The northern teams enter the city sector, with the mission of clearing two separate buildings (Buildings 1 and 2) and then continue to secure the Objective, a building in the center of the area hosting an important opposition meeting. The southwestern teams enter the sector to secure Buildings 3 and 4 and then continue to the Objective. The western teams (HQ attachment) are tasked to proceed directly to the Objective. The interior resistance is provided by opposition DI located in the five critical buildings and also located in a key vantage point building on the northwest of the sector. Three additional opposition DI are stationed outside of Buildings 1, 2 and the Objective.

A version of the One Semi-Automated Forces (OneSAF) simulation, known as DISAF v9.4, serves as the medium for scenario development and a detailed rendition of close fighting conditions among individual combatants. Data collected from 300 scenario runs supports the parameter investigation. Each run ends with the Objective taken or the attacking fireteams killed.

3. BATTLE PARAMETERS

An extensive set of battle parameters and response measures are possible in a simulated environment. The posture taken in this research is that any number of new measures not currently available in practice may become so as network centric warfare becomes more a reality. Thus, a mixture of traditional measures (e.g., mobility kill, catastrophic kill) and currently unavailable measures (e.g., real-time round expenditure, number of hits taken) are assumed available in this study.

Efficient data collection was accomplished by inserting into OneSAF a killer/victim scoreboard code to collect time-stamped entity fire exchanges and logistics data (Heilman and O'May, 2002). We also utilized shell scripts to parse the data into rollups of different aspects of the battle (e.g., the amount of each type of ammunition expended).

Information on battlefield entities was expressed in terms of engagement activity and status for varying entity aggregations and at different times. Engagement activity is defined as the number of shots fired and shots received and is recorded uniquely for each pair (shooter, target). Status refers to the health of the entity and is expressed differently according to type: a soldier was either dead or alive; a vehicle was undamaged or designated with a gradated kill (mobility, firepower, mobility and firepower, or catastrophic).

DI engagement data from the simulation was generally aggregated according to fireteam or building, but was individually scored for the three opposition DI stationed outside Buildings 1, 2, and the Objective. Fireteam status was aggregated in the same manner, with a value awarded according to the remaining firepower of the team (SAW[3], M203[2], Lfk5[2], Ags17[2], AK47[1], M16[1]). Engagement activity and status were similarly grouped by type and location for vehicles, with undamaged vehicles considered to have a value of four. Degraded vehicles received a reduced value according to kill (mobility[3], firepower[2], mobility and firepower[1], catastrophic[0]).

Engagement information was reported for all relevant pairings of shooter and target. Some pairings were not relevant because the entities never engaged. Status and engagement information combined required 298 parameters.

Initial summaries were reported for two fixed times (time slices) during the battle. The first, ~ 5.5 minutes into the battle, represented the approximate time it took the DI to reach the sector. The second, ~ 8.0 minutes into the battle, corresponded to a time approximately half-way through the average battle. These two times were characterized by 596 parameters.

Subsequent to the initial summaries, the data was augmented by further aggregation in two steps. In the first step, forces were partitioned according to approach (north, west, southwest) or defensive position (north, west, southwest, or center). Status values for these seven groups were provided for each fixed time, adding fourteen additional measures. In the second step, this augmentation was carried still further, providing status values for the seven groups at each minute mark beginning at 1 minute into the battle and ending at 11 minutes into the battle, excluding status values at 8 minutes, which corresponded to the second fixed time. The additional seventy values inflated the number of parameters for consideration to 680.

Seven measures reflected aspects of battle outcome. Five signified the taking of individual Buildings 1–4 and the Objective. An additional measure (MOUTscore) cumulatively assessed points for ground taken, while discounting that total according to remaining attacking force strength. The Objective was assessed 4 points; Buildings 1–4 were each assessed 1 point. The point total for the ground taken was halved if fewer than 12 friendly DI survived. A final measure, termed Foothold, sought to represent ground taken by direction. Unlike MOUTscore,

Foothold is a categorical measure, but with an implied ranking, assigning a code according to having secured no buildings (1) the north or south buildings (2), both the north and south buildings (3), the Objective (4), or all buildings (5).

4. ANALYSIS

The data mining goal of the analysis is to uncover useful structure in the data, which in turn provides new insight or clarifies understanding of the relationship between battle parameters and battle outcome. We approach this in three steps. First, we explore the suitability of available outcome measures to serve as a response in the analysis. Second, we reduce the dimensionality of the problem with phased parameter screening to support more formal modeling. Finally, we establish candidate models relating the remaining parameters to the most informative outcome measures.

4.1 Outcome Summary

Outcomes measures discussed in Section 3 include five binary measures, one score, and one categorical measure. Binary measures indicate only whether the component was taken at the end of the battle. In Table 1, we see the results for the 300 scenario runs. Of course, in practice we would always like to take the Objective and intermediate building objectives, but in this simulation exercise, we learn about the relationship between battle parameters and outcome only if the outcome is varied. Consequently, the intermediate objective for the southwest fireteams in taking Building 3 will probably yield little information related to battle parameters. Each of the other binary outcomes shows sufficient variability to be of some use directly or indirectly through the MOUTscore and Foothold.

Table 1. Outcome summary for binary measures

Response	Success (1)	Failure (0)	Rate
Objective	196	104	65.3%
Bldg 1	208	92	69.3%
Bldg 2	211	89	70.3%
Bldg 3	292	8	97.3%
Bldg 4	168	132	56.0%

The distribution of MOUTscore appears in Figure 2. Not surprisingly, given the 2/3-success rate for attacking forces, higher MOUTscores tend to occur more frequently. Note there is variability in the score that may be explained by specific battle parameters.

The distribution of Foothold appears in Figure 3. This measure shows considerable variability among its possible classes. Its principal advantage over the MOUTscore is a more natural grouping of intermediate objectives and lack of dependence on relative point values for intermediate and final objectives.

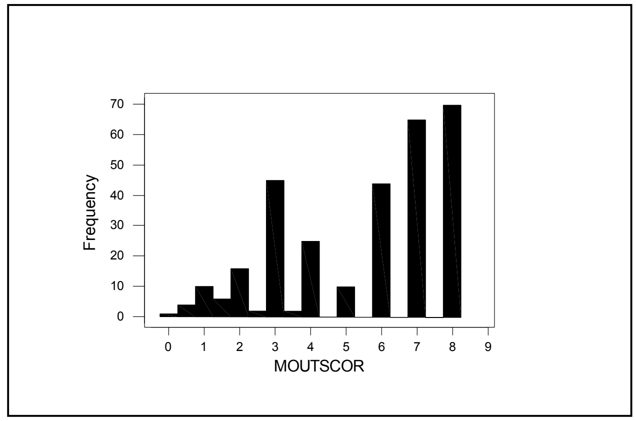

Fig. 2. Frequency distribution of MOUTscore

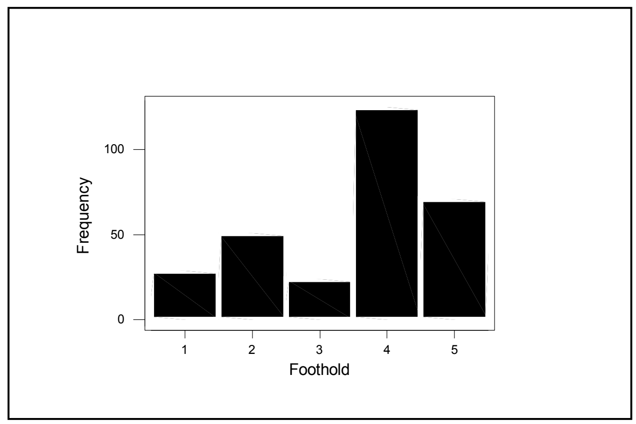

Fig. 3. Frequency distribution of Foothold

The elapsed time to reach an outcome reflects on the importance of time-dependent battle parameters introduced in Section 4.2. The median elapsed time was ~ 16 minutes with the inner 50% of values ranging between ~ 15 and ~ 18 minutes. The distribution was skewed to the right with a minimum of ~ 12 minutes and a maximum of ~ 45 minutes.

4.2 Parameter Screening

Through the killer/victim scoreboard we can collect large amounts of data, but there is no guarantee of their utility. The most basic requirement for a candidate explanatory measure is that it varies; without variation there is no basis for association with a response measure. Among the 680 potential explanatory measures, 289 failed to vary, leaving 391 variables to explore.

The remaining variables number too many to process through most available statistical routines. A second phase

of screening allows us to determine prediction potential. To accomplish this, we performed separate two-sample t-tests for each of the 391 variables, where samples arose from the two conditions of taking or not taking the Objective. The t-test is not to be regarded as a formal instrument in this situation but merely as a means to gauge belief that separation of variable values occurs according to battle outcome. We then ordered the variables according to their p-value in the statistical test and arbitrarily retained for further consideration the 91 measures having a p-value less than 0.05. This rough ranking procedure was in many cases supported with graphical inspection. Figure 4 associates the number of shots taken by M2s approaching from the southwest directed toward the T80 defending that approach. The jittered (random offset) values along both axes allow us to see the density of observations at each pairing and clearly associate taking the objective with more fire laid on this T80. Similarly, the same measure is associated with Foothold in a sensible fashion in Figure 5. If we accept the implied ranking of the Foothold measure, higher values suggest a better outcome.

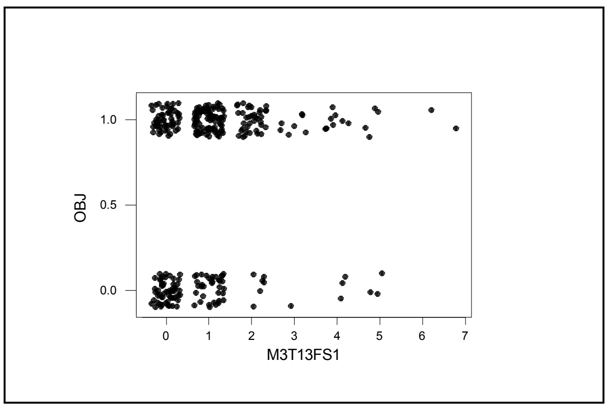

Fig 4. Shots fired by M2s approaching from the southwest directed toward a defending T80, associated with whether or not the objective is taken

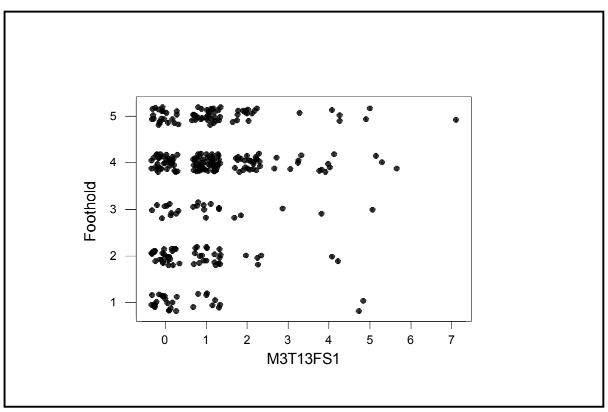

Fig 5. Shots fired by M2s approaching from the southwest directed toward a defending T80, associated with Foothold

4.3 Modeling

In this section we present excerpts from the complete modeling effort to demonstrate the potential for associating battle parameters with battle outcome. The two approaches entertained are statistical discriminant analysis and classification and regression trees.

A third screening phase merely ensures that no variables remaining are perfectly correlated. High correlation among explanatory variables causes ill-conditioned matrices and poor model development for many statistical procedures. Perfect correlation occurred in some instances when the same measure was taken at different times. If no change occurred, a perfect correlation resulted. In those instances, we used only the variable corresponding to the earlier time.

Analytical details for modeling are not pursued here, but we do point out that in such modeling the analyst should strive to limit the confusion in the classification matrix and explain as much variation of the response with as few variables as possible.

4.3.1 Example 1

In this first example, a stepwise discriminant analysis was performed using as a response, taking the mission objective. The classification results for the learning matrix are given in Table 2. The last two columns indicate predicted class membership where "p" indicates the actual proportion of the data falling in that class. We see that there is considerable ambiguity in recognizing early an unfavorable outcome as we correctly did so only ~ 56% of the time. However, successful outcomes are recognized in advance ~ 90% of the time.

Table 2. Initial stepwise discriminant classification

Classification Matrix			
Observed Class	Percent Correct	0 p=.3467	1 p=.6533
0	55.77	58	46
1	90.31	19	177
Total	78.33	77	223

The classification results for the learning sample provide some encouragement that parameters used in the model possess relevant information regarding the response. Nine of the 91 measures appear in this default model. We discuss some of these. Parameters F5TS1 and F3TS2 are remaining strength measures for Fireteams 5 and 3 taken at time slices 1 and 2, respectively. Fireteam 5 approaches from the west and Fireteam 3 approaches from the north. Their success in attacking positions to the north appears to be especially important in completing the mission.

Parameter F3A21FS1 counts the number of shots directed by Fireteam 3 from the north toward a key opposing DI. The parameter suggests that as more shots are required, the chance for eventual success diminishes. This is plausible because Fireteam 3 should be the last of three northern fireteams to engage this heavily armed DI entity. It suggests that the Fireteams 1 and 2 did not fair well in the northern approach, which would tend to reduce the chance of overall success.

Similarly, T1T11FS1 records the number of shots directed by an M1A1 approaching from the west toward a T80 in a defensive posture between Buildings 1 and 2 to the north. If shots must be taken by the M1A1, the chance for capturing the Objective decreases. Further evaluation shows that in ~ 80% of the battles where the M1A1 engaged the T80, at least one of the M2s from the north had lost mobility and firepower, suggesting stiff resistance from the northern defense.

More intuitive is the role of parameter F6BL5FS1, which tallies the shots in an engagement by Fireteam 6 from the southwest directed to the northwest position of Building 5. This would not occur unless the southwest defense was sufficiently softened to allow the advance of Fireteam 6. This appears to be supported by the observation that the amount of fire placed on Building 5 by Fireteam 6 increases to 2:1 when the T80 defending the southwest approach is firepower killed (F) compared with undamaged (U) (Table 3).

Additionally, combined strength measures for the southwest approach and the center of town defense recorded at eleven minutes into the battle were included in the classification function. Even at eleven minutes, these measures are valuable, because on average ~ 68% of the blue force strength remains for this scenario.

Table 3. Frequency of "n" shots fired by T80 status

T80 State	Shots Fired (n) by Fireteam 6 on Building 5						
	0	1	2	3	4	5	6
F	67	52	44	25	12	2	1
U	34	22	25	6	3	0	0

4.3.2 Example 2

Our second example illustrates the use of classification and regression trees. Taking the Objective remains our response, but we allow a version of classification and regression trees from StatSoft© to perform variable selection of the parameters. Prior to the variable selection, we extended our parameter screening beyond what we had done for the discriminant analysis in Example 1. The version of software we used would only consider a maximum of fifty parameters. We selected the best fifty using the p-values referenced previously. The classification matrix for this alternative approach appears as Table 4. Note the overall classification performance is slightly worse than that of discriminant analysis. We recognize early an unfavorable outcome only ~ 53% of the time and favorable outcomes are recognized in advance ~ 85% of the time. However, this model does have some advantages in that it relies on only five parameters instead of nine and that, generally, results are easier to interpret.

Table 4. Classification and regression trees

Classification Matrix			
Observed Class	Percent Correct	0 p=.3467	1 p=.6533
0	52.88	55	49
1	84.69	30	166
Total	74.33	77	223

The approach of classification and regression trees is favored by many because of a more intuitive interpretation of parameters included in the model. Figure 6 displays the classification tree corresponding to Table 4. In this tree, there are five parameters: BSWTS11, F3TS2, F4BL1FS2, RCTS11, and BMB11TS2 used to determine the classification of each scenario run. Six terminal nodes (moving down) for the tree result from splits on the five parameters. The path to each terminal node identifies the conditions for classifying a run as a success (1) or failure (0). The membership of each node in the tree is given in Table 5. The class assignment appears in the upper right hand corner of each node.

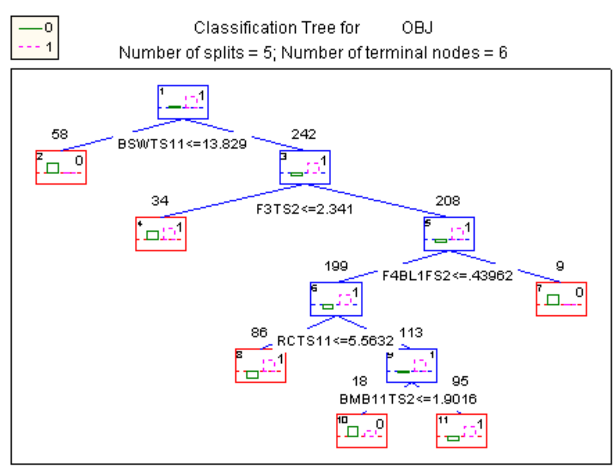

Fig 6. Classification tree

Interpretation proceeds sequentially. We will describe the first two steps. The first split is on the parameter BSWTS11, the aggregate strength of the southwest approach eleven minutes into the battle. If this strengths falls below 13.8 on an initial strength of 27 (~ 51%) the

run is immediately classified as a failure. Those scenario runs are sent to the left in the tree. Of the 58 runs classified as a failure, 39 (~ 67%) actually are failures as indicated in Table 5. The remaining 242 runs are temporarily classified as successes. The importance of BSWTS11 is not surprising. On many occasions, southwest forces took the Objective, sometimes sweeping to the east of building 3 before proceeding to the Objective. Next the parameter F3TS2, the strength of Fireteam 3 eight minutes into battle is considered. If F3TS2 has a value of 2 or less, the DI with the SAW has been lost along with at least one DI carrying an M16. The importance of this northern Fireteam was mentioned in Example 1. A weakened Fireteam 3 causes the run to be classified a failure. Note in Table 5 this classification is only ~ 50% accurate.

The only engagement activity entering the model is the parameter F4BL1FS2, which contains information on the fire laid on Building 1 by Fireteam 4 approaching from the west. The model suggests less is more with regard to reaching the Objective. This again makes sense because if Fireteam 4 has to engage Building 1, it is because Fireteams 1–3 have not completed the job.

The process continues with additional remaining strength measures for the defending forces in the neighborhood of the Objective and for two attacking force M2s.With a little practice, the impact of parameters in classification tree models is easily established.

Table 5. Predicted classes and split conditions

Tree Structure (FINALMOUTEX9_04)
Child nodes, observed class n's,
predicted class, and split condition for each node

Node	Left branch	Right branch	n in cls 0	n in cls 1	Predict. class	Split constant	Split variable
1	2	3	104	196	1	-13.8287	BSWTS11
2			39	19	0		
3	4	5	65	177	1	-2.3410	F3TS2
4			16	18	1		
5	6	7	49	159	1	-0.4396	F4BL1FS2
6	8	9	43	156	1	-5.5632	RCTS11
7			6	3	0		
8			10	76	1		
9	10	11	33	80	1	-1.9016	BMB11TS2
10			10	8	0		
11			23	72	1		

4.3.3 Example 3

We attempted to model the more complicated measure, Foothold, using discriminant analysis as well. The results are less than satisfying, in part because of inherent difficulty in accurately classifying five groups, but for completeness we have included the classification matrix from the software as Table 6. Immediately, we see the overall ability to classify correctly the Foothold level achieved is only ~ 56%. A bright spot is the ability to predict the level 4 Foothold ~ 81% of the time. However, given the definition of the Foothold score, this result is merely an affirmation of the Objective as a predictable outcome. Additional work is required to develop a more informative measure that gives a gradation of performance while remaining accessible to standard modeling approaches.

Table 6. Classification matrix for Foothold parameter

Classification Matrix (FINALMOUTEX9_04)
Rows: Observed classifications
Columns: Predicted classifications

Class	Percent Correct	1 p=.0967	2 p=.1700	3 p=.0800	4 p=.4167	5 p=.2367
1	48.27586	14.00000	4.00000	1.000000	10.0000	0.00000
2	33.33333	3.00000	17.00000	0.000000	23.0000	8.00000
3	16.66667	0.00000	4.00000	4.000000	8.0000	8.00000
4	80.80000	2.00000	7.00000	0.000000	101.0000	15.00000
5	46.47887	0.00000	3.00000	3.000000	32.0000	33.00000
Total	56.33333	19.00000	35.00000	8.000000	174.0000	64.00000

Classification Matrix (FINALMOUTEX9_04)

CONCLUSIONS

In this exercise we have demonstrated a capability to describe in intricate detail dynamic battle information—and to learn from it. The principal finding of this work is that data mining of urban combat simulation holds promise for understanding in fidelity previously not possible the importance of specific units and their spatial and temporal characteristics. The implications of such a capability are far-reaching for the warfighter. With such information, course of action development for a planned assault could be better evaluated and perhaps refined. Key measures during battle execution might be uncovered through the data mining process, providing markers along the way for the commander to gauge the progress of his forces.

The process however remains difficult, especially so for this urban environment. Problems remain with the simulation used to model the engagement. It takes considerable time to develop the scenario and to execute it a sufficient number of times to support data mining. Analysis is far from automated at this point and requires someone knowledgeable in statistics or data mining to reveal the important data structure. Measures for mission success and to describe entity characteristics remain elusive. Although useful information is gleaned with the measures we now have, we feel that any improvements would greatly enhance the analysis. One especially difficult issue is the granularity required to adequately describe the state space. We experimented with different levels of aggregation but still struggle with porridge too hot or too cold.

Future work must focus on these nagging problems. We now have identified software to facilitate automation of the analysis and intend to work toward independence from the analyst. Though outside our work area, newer versions of OneSAF and other simulations promise to improve the behaviors and reliability of combat entities. In addition, more work is intended utilizing high performance computing resources to speed individual scenario runs and reduce the time required to build the data set. We have experimented with this for limited runs, but have not leveraged that tool to the fullest. We will continue to refine our methods through further experimentation to examine combat on urban and complex terrain.

REFERENCES

Bodt, B., Forester, J., Hansen, C., Heilman, E., Kaste, R., and O'May, J., 2002: Data Mining Combat Simulations: A New Approach to Battlefield Parameterization, *Proceedings of the 23rd Army Science Conference*, Orlando, FL.

Heilman, E. and O'May, J., 2002: One Semi-Automated Forces (OneSAF) Killer/Victim Scoreboard (KVS) Capability, *Army Research Laboratory Technical Report, ARL-TR-2829*, Aberdeen Proving Ground, MD.

Liao, T., Bodt, B., Forester, J., Hansen, C., Heilman, E., Kaste, R., and O'May, J., 2003: Discovery of Battle States Knowledge from Multi-Dimensional Time Series Data, *Proceedings of the 35th Symposium on the Interface: Computing Science and Statistics*, Salt Lake City, UT.